

Remarks

Claims 1-42 are currently pending in this application.

Objections to the Drawings

5 The Examiner objects to the drawings under 37 CFR 1.83(a) as not showing every feature of the invention specified in the claims. Specifically, the Examiner suggests that the “hash table” and the “tree data structure” must be shown. The Applicant traverses this objection.

 These elements are represented in FIG. 1 by, for example, Answer Cache 165 and
10 Referral Cache 168. These elements are disclosed to include a hash table and tree data structure, according to some embodiments, in paragraphs [0026] and [0027] of the specification as filed. 37 CFR 1.83(a) requires that where a “detailed illustration is not essential for a proper understanding of the invention, [features] should be illustrated in the drawing in the form of a graphical drawing symbol or a labeled representation (e.g., a
15 labeled rectangular box).” In the current application, a detailed illustration is not essential for a proper understanding of the invention because hash tables and tree data structures, generally, are well known in the art. Because these elements are well known in the art, the applicant does not have to provide an illustration that describes further detail of these elements (See MPEP 2106.II.C).

20

Claim Objections

Claims 10 and 18 are objected to on the ground that “an Mx record” should read “a Mx record.” The Applicant appreciates the Examiner’s suggestion for this modification. However, the Applicant believes that the claim language is correct and

wishes to avoid making any unnecessary claim amendments that might be improperly interpreted as being required for allowance. As written, Claims 10 and 18 meet the requirements of §112.

Claim 25 is objected to under 37 CFR 1.75(c) as being of improper dependent
5 form for failing to further limit the subject matter of previous Claim 19. The Applicant traverses this objection.

Claim 19 recites, in part, “*a hash table configured to store the answer information or to store a pointer to the answer information.*” The “or” separates two alternatives, “*to store the answer information,*” and “*to store a pointer to the answer information.*” In
10 Claim 19, the first alternative is optional if the second alternative is satisfied. Claim 25 recites, “*wherein the hash table is configured to store the answer information,*” and thus requires that the hash table is configured to store the answer information. Claim 25, thus, is narrower in scope than Claim 19 because the first alternative is no longer optional. It is, therefore, the position of the Applicant that Claim 25 meets the requirements of 37
15 CFR 1.75(c).

Rejections under 35 USC §102

Claims 1-9 and 11-42 are rejected under 35 U.S.C. 102(e) as being anticipated by Hudson Michel (2002/0143984) hereinafter, Hudson.

Regarding Claim 1,

20 **Claim 1 recites:**

- 1. A caching server comprising:
an answer cache configured to access answer information through a flat data structure;
a referral cache configured to store referral information; and*

computer instructions configured to translate a domain name into DNS information by examining the answer cache and, responsive to the results of examining the answer cache, examining the referral cache.

5 In rejecting Claim 1, the Examiner states “Hudson teaches a caching server (i.e. the web cache server) comprising an answer cache (a proximal cooperative cache) configured to access answer information through a flat data structure; a referral cache (a distal cooperative cache) configured to store referral information; ...” The Applicant traverses this statement.

10 First, the cache servers taught by Hudson are taught to cache “web content data.” For example, paragraph [0004] of Hudson characterizes “web content data” as web pages and images. Thus, the caches of Hudson are traditional web page caches configured to speed the delivery of web pages to a requester. This interpretation of the teachings of Hudson is supported by paragraphs [0003] – [0006] wherein Hudson teaches “a web
15 caching system consists of one or more caches that store copies of web pages, images and other web content data.” The web page caches of Hudson are not equivalent to the answer cache and referral cache of the present application.

In construing the meaning of the terms “*answer cache*” and “*referral cache*,” “The broadest reasonable interpretation of the claims must ... be consistent with the
20 interpretation that those skilled in the art would reach,” (MPEP 2111), and “[o]ffice personnel must rely on the applicant’s disclosure to properly determine the meaning of the claims, (Markmen v. Westview Instruments, MPEP 2106.II.C). The answer cache and referral cache, of Claim 1, are specific types of DNS (Domain Name Service) caches configured for the resolution of domain names into IP (Internet Protocol) addresses.
25 These are not “web content data” caches. As would be understood by a person of

ordinary skill in the art, the answer cache is configured to store IP addresses in association with domain names and to provide an IP address in response to a domain name, and the referral cache is configured to refer a request for an IP address to an answer cache. These definitions are supported in paragraphs [0023], [0026] and [0027] of the application as filed. The answer and referral caches of the claims are, thus, not the same as the web page caches of the cited art. To equate the web page caches of the cited art with the “*answer cache*” and “*referral cache*” of Claim 1, would be inconsistent with both the interpretation that those skilled in the art would reach and the disclosure as filed.

In pointing out these differences between answer and referral caches and web page caches, the Applicant is not attempting to read limitations into the claims. Rather, the Applicant is pointing out how a person of ordinary skill in the art would interpret these terms. A person of ordinary skill in the art would appreciate substantial distinctions between a web page cache and answer and referral caches. Thus, a person of ordinary skill in the art would not equate the web page caches of Hudson with the answer and referral caches of Claim 1.

Second, it is the position of the Applicant that Hudson teaches the creation of a forwarding table using a hash code, but does not teach that the forwarding table itself has a flat data structure. Thus, the use of the forwarding table to access information does not include the use of a flat data structure, and Hudson does not teach, “*an answer cache configured to access answer information through a flat data structure,*” as recited in Claim 1.

Specifically, at paragraph [0018], Hudson teaches, “creating forwarding table using hash function encoding.” Examples of hash codes used by these functions are

provided in FIGs. 4A and 4B, ([0036] and [0037]) of Hudson. For example, “org” is taught to encode as “00074bea.” Once the appropriate encoding is found using the hash table, the code is stored in a forwarding table. Hudson specifically teaches that this storage includes a tree data structure (e.g., **not** a flat data structure). For example, in the
5 second half of paragraph [0031], Hudson teaches “[t]he sequence of hash codes are used ... to locate the decomposed URL in the URL decomposition **tree** in the forwarding table.” This URL decomposition tree is further defined in paragraph [0027].

While the tables illustrated in FIGs. 4A and 4B of Hudson may be considered flat data structures, Hudson does not teach an “*answer cache configured to access answer*
10 *information through*” these tables. The FIGs. 4A and 4B of Hudson are merely used to generate codes to be stored in a tree data structure of a cache. Thus, Hudson does not teach a cache having a “*flat data structure*,” and does not access to “*information through a flat data structure*,” as recited in Claim 1. The Applicant, therefore, requests that the Examiner more specifically explain how Hudson is thought to teach, “*an answer cache*
15 *configured to access answer information through a flat data structure*,” as recited in Claim 1, or allow Claim 1.

Third, it is the position of the Applicant that Hudson teaches the use of a hash table in developing a referral cache not an answer cache. Therefore, Hudson does not teach, “*an answer cache configured to access answer information through a flat data*
20 *structure*,” as recited in Claim 1. Thus, even if one were to assume for the sake of argument that the teachings cited by the Examiner include a flat data structure, Hudson would not teach the limitations of Claim 1.

Specifically, the forwarding table of Hudson is defined as being “used to translate received web request URLs into IPAs for forwarding web requests to the cooperative caches storing the requested web content data,” (Paragraph [0056]). Thus, even assuming for the sake of argument that the forwarding table has a flat data structure, this definition more closely matches that of a referral cache rather than an answer cache. See, for example, paragraphs [0022] and [0023] of Hudson. The Applicant, therefore, requests that the Examiner point out teaching of “*a flat data structure*” for “*an answer cache*” to access information, or allow Claim 1.

Fourth, it is the position of the Applicant that Hudson does not teach the “*answer information*,” as recited in Claim 1. In rejecting Claim 1, the Examiner infers this information is stored in “a proximal cooperative cache” and “a distal cooperative cache.” However, Hudson teaches that these caches are used to store “web content data,” (see paragraphs [0003] and [0005]). The Applicant understands “web content data” as meaning a webpage or the like, referred to by a URL (see paragraph [0004] of Hudson).

In contrast, the specification as filed defines “answer information” as follows: “[t]he answer cache is configured to store answer information such as domain names and associated IP addresses.” (See paragraph [0022]). This definition conforms to the meaning that a person of ordinary skill in the art would have for “answer information.” The Applicant is unable to identify any teaching within Hudson that URL web content data includes “*domain names and associated IP addresses*.” In pointing out these differences between an answer cache and a web content cache, the Applicant is not attempting to read limitations into the claims. Rather, the Applicant is pointing out how a person of ordinary skill in the art would interpret these claim terms. As pointed out

above, “The broadest reasonable interpretation of the claims must ... be consistent with the interpretation that those skilled in the art would reach,” (MPEP 2111), and “[o]ffice personnel must rely on the applicant’s disclosure to properly determine the meaning of the claims, (Markmen v. Westview Instruments, MPEP 2106.II.C). To interpret “*answer*”
5 *information*” as including URL web content data would be inconsistent with both the interpretation that those skilled in the art would reach and the specification as filed.

For at least the above reasons, the Applicant requests that the Examiner more specifically point out these limitations of Claim 1, or allow Claim 1 and those claims that depend therefrom.

10 **Regarding Claims 11, 13, 17, 26, 34 and 38-39,**

The Examiner rejects Claims 11, 13, 17, 26, 34 and 38-39 on the same rationale as Claim 1. It is the position of the Applicant that Claims 11, 13, 17, 26, 34 and 38-39 are allowable for at least the reasons discussed above with respect to Claim 1. Further:

Claim 11 recites:

- 15 11. *A computer readable medium having stored thereupon computer code configured to determine DNS information associated with a domain name, the computer code comprising:*
a code segment configured to receive a request for the DNS information
corresponding to a domain name;
20 *a code segment configured to examine a first cache to find the DNS information, the first cache including a flat data structure and configured to store the DNS information or a pointer to the DNS information; and*
a code segment configured to initiate a search of a second cache if the DNS
25 *information is not found by examining the first cache, the second cache configured to store data referring to further locations on a computer network wherein the DNS information may be found.*

The Applicant traverses the rejection of Claim 11 on the grounds that Hudson teaches requests for “URL web content data,” while Claim 11 recites “*a request for the*

DNS information.” Paragraph [0004] of Hudson characterizes “web content data” as including “web pages” and “images.” “Web pages” and “images” are not “*DNS information corresponding to a domain name.*” Thus, the requests for “URL web content data” of Hudson do not teach “*a request for the DNS information,*” as recited in Claim 11. The Applicant, therefore requests that the Examiner specifically point out teachings of “*a code segment configured to receive a request for the DNS information corresponding to a domain name,*” or allow Claim 11, and those claims that depend therefrom.

Further, Claim 11 recites:

*a code segment configured to examine a first cache to find the DNS information, the first cache including a flat data structure and configured to store the DNS information or a pointer to the DNS information; and
a code segment configured to initiate a search of a second cache if the DNS information is not found by examining the first cache, the second cache configured to store data referring to further locations on a computer network wherein the DNS information may be found.*

Neither of these limitations are included in Claim 1 and, thus, do not appear to be addressed by the Examiner. The Applicant requests that the Examiner specifically point out all the limitations of Claim 11, within the cited art, or allow Claim 11 and those claims that depend therefrom.

Regarding Claims 13, 17 and 26, the Applicant believes that Claims 13, 17 and 26 (and those claims that depend therefrom) are allowable for at least the same reasons as Claims 1 and 11.

Regarding Claim 34 and 38-39, the Applicant believes that Claims 34 and 38-39 are allowable for at least the same reasons as Claims 1, 11, and Claim 33 from which

they depend. The Applicant also believes that Claim 39 is allowable for the reasons discussed herein with respect to Claim 19.

Regarding Claim 2, the Applicant believes that Claim 2 is allowable for at least the reasons discussed with respect to Claim 1.

5 **Regarding Claims 25, 29 and 35,**

Claim 25 recites:

25. The method of claim 19, wherein the hash table is configured to store the answer information.

10 The Examiner rejects Claim 25 on the same rationale as Claim 2. The Applicant traverses this rejection. Specifically, paragraph [0016] and FIG. 4A of Hudson, as cited by the Examiner in the rejection of Claim 2, do not include answer information. Both the cited text and figure concern encoding information for compressed storage of URLs, but they do not include “*answer information*” as recited in Claim 25. As discussed above in
15 regard to Claim 1, “*answer information*” would be interpreted by one of ordinary skill in the art as including “*domain names and associated IP addresses.*” In contrast, as illustrated in FIG. 3, the encoding information of Hudson includes parts of a URL to be stored in a tree data structure (of FIG. 3). The text cited by the Examiner does not teach the limitations of Claim 25 because the encoded information of FIG. 4A does not include
20 IP addresses and also because the encoded information is stored in a tree data structure rather than a hash table. The Applicant, therefore, requests that the Examiner specifically point out how the cited teachings of Hudson teach, “*answer information*” or allow Claim 25.

The Applicant further believes that Claim 25 is allowable for the reasons
25 discussed herein with regard to Claim 19, from which it depends.

Regarding Claims 29 and 35, the Applicant believes that Claims 29 and 35 are allowable for at least the reasons discussed with respect to Claims 1 and 2, and Claims 26 and 33 from which they depend.

Regarding Claim 19,

5 **Claim 19 recites:**

10 19. *A method of determining DNS information, the method comprising:
receiving a request for DNS information corresponding to a domain name;
examining an answer cache for answer information, the answer cache including a
hash table configured to store the answer information or to store a pointer
to the answer information; and
searching a tree data structure if the DNS information is not found by examining
the answer cache.*

The Examiner rejects Claim 19 on the same rationale as the rejection of Claims 1
15 and 2. However, it is the Applicant's position that the Examiner has not pointed out
teaching of an "*answer cache including a hash table*" in combination with "*searching a
tree data structure if the DNS information is not found by examining the answer cache,*"
as recited in Claim 19. The cited art does not appear to teach that one cache has a
different data structure than the other cache, much less that one cache includes a hash
20 table while the other cache includes a tree data structure. The Applicant, therefore,
requests that the Examiner specifically point out these limitations, or allow Claim 19, and
those claims that depend therefrom.

Further, the Applicant believes that Claim 19 is allowable for at least the reasons
discussed herein with respect to Claims 1, 11 and 22.

25 **Regarding Claim 3,**

Claim 3 recites:

3. *The caching server of claim 1, wherein the flat data structure includes pointers to a
tree data structure.*

In rejecting Claim 3, the Examiner states, "Hudson teaches the claimed invention as described above and furthermore, Hudson teaches that the flat data structure includes pointers to a tree data structure (shown in Fig. 5)." The Applicant traverses this
5 statement.

FIG. 5 is labeled "Universal Resource Locator Forwarding Table." It is the Applicant's understanding that this is the same "forwarding table" that the Examiner has argued has a flat data structure in the rejection of Claim 1, and other claims. The forwarding table of FIG. 5 includes pointers from elements of a tree data structure to
10 other elements of the same tree data structure. Further, a tree data structure is not a flat data structure. Thus, a teaching of pointers between elements of a tree data structure does not teach a "*flat data structure include[ing] pointers to a tree data structure,*" as recited in Claim 3.

Further, the Applicant respectfully points out that the Examiners' arguments with
15 respect to Claims 1 and 3 are contradictory. In one instance, it is argued that the forwarding table of FIG. 5 teaches a flat data structure, and in the other, it is argued that the same forwarding table includes pointers from a flat data structure to a tree data structure. If the forwarding table of Hudson has a tree data structure, rather than a flat data structure as suggested by the Examiner in the rejection of Claim 1, then Claim 1,
20 among others, should be allowable. Alternatively, if the forwarding table of Hudson does not have a tree data structure, then Claim 3 should be allowable. As stated above, it is the position of the Applicant that the forwarding table of Hudson has a tree data structure, thus Claim 1 (and others) are allowable, and Claim 3 is allowable, at least, as being dependent on Claim 1.

Regarding Claims 20-23 and 27-28,

In rejecting Claims 20-23 and 27-28, the Examiner states “see arguments with respect to the rejection of claim 3. Claims 20-23 and 27-28 are also rejected based on the same rationale as claim 3.” The Applicant traverses these rejections.

5 First, it is the Applicant’s position that Claims 20-23 and 27-28 are allowable for at least the reasons discussed above with respect to Claims 1 and 3.

Furthermore:

Claim 20 recites:

10 20. *The method of claim 19, wherein the hash table is configured to store the pointer to the answer information.*

While FIG. 5 of Hudson illustrates the use of pointers, it is unclear to the Applicant how any of these pointers are taught to point to “*answer information*,” as recited in Claim 20. The pointers within FIG. 5 appear to point to fragments of a
15 disassembled URL within the same tree data structure. The Examiner is referred to the examples of “*answer information*” discussed above with respect to Claims 1 and 25, which show that a disassembled URL is not “*answer information*.” The Applicant, therefore, requests that the Examiner specifically point out which aspects of the cited art are believed to teach “*answer information*” and how FIG. 5 teaches pointers from a hash
20 table to “*answer information*,” or allow Claim 20.

Claim 21 recites:

21. *The method of claim 19, wherein the answer cache does not include a tree data structure.*

25 As with the rejection of Claim 1, the Applicant finds the Examiner’s rationale for rejection of Claim 21 contradictory to the rejection of Claim 3. The Examiner has

suggested that the forwarding table of Hudson teaches the answer cache having a flat data structure (discussing Claim 1) and that the forwarding table shown FIG. 5 of Hudson teaches pointers to a tree data structure (discussing Claim 3). This implies that the forwarding table assumed to teach the answer cache includes a tree data structure.

5 However, Claim 21 specifically recites, “*the answer cache does not include a tree data structure.*” As such, the rejection of Claim 21 is contradictory to the rejection of Claim 3.

Claim 22 recites:

10 22. *The method of claim 19, wherein the tree data structure is configured to store referral data and is included in a referral cache.*

Claims 19 and 22 recite searching within a referral cache including the tree data structure “*if the DNS information is not found by examining the answer cache.*” This would not be possible under the Examiner’s interpretation of Claims 1 and 9. In the
15 Examiner’s view of Hudson, the DNS information includes pointers to the referral cache. If these pointers were not found then the referral cache could not be found and searching within a referral cache including the tree data structure “*if the DNS information is not found by examining the answer cache,*” would be impossible.

Specifically, according to the Examiner’s rejection of Claim 9, the “*DNS information*” includes the IP addresses used to forward a request to a proximal web
20 content data cache, and according to the Examiner’s rejection of Claim 1 the “*answer cache*” is taught by the proximal web content data cache. Assuming, for the sake of argument, that these interpretations were correct, then the IP addresses would be required to forward a request and if the suggested DNS information (IP addresses) were not found,
25 then the suggested answer cache (web content data cache) could not be found because the

IP addresses would be unavailable. Thus, the cited art cannot teach searching a referral cache “*if the DNS information is not found,*” as recited in Claim 22.

The Applicant further believes that Claim 22 is allowable for at least the reasons discussed herein with respect to Claim 19, from which it depends.

5 **Regarding Claims 23, 27 and 28**, the Applicant believes that Claims 23, 27 and 28 are allowable for at least the reasons discussed with respect to Claims 19 and 26, from which they depend.

Regarding Claim 4, the Applicant believes that Claim 4 is allowable for at least the reasons discussed herein with respect to Claims 1, 3 and 19. Specifically, the cited art
10 does not include a tree data structure configured to store answer information and does not include a flat data structure configured for accessing answer information.

Regarding Claim 5, the Applicant believes that Claim 5 is allowable for at least the reasons discussed herein with respect to Claims 1, 3 and 19. Specifically, the cited art does not include a flat data structure within an answer cache, the flat data structure
15 including pointers to a tree data structure within a referral cache.

Regarding Claim 6,

Claim 6 recites:

20 6. *The caching server of claim 1, wherein the caching server is also an authoritative server.*

 With regard to Claim 6, the Examiner states “Hudson teaches that the caching server (i.e. the web cache server) is also an authoritative server, i.e. a server which has the desired information (e.g. see the abstract).” The Applicant traverses the Examiner’s definition of “authoritative server.” It is the position of the Applicant that a person of
25 ordinary skill in the art would not define “authoritative server” as “a server which has the

desired information,” as suggested by the Examiner. In the context of DNS services, the term authoritative server has a more specific meaning. The Examiner is referred to paragraph [0004] of the specification as filed which states “[i]nformation to be looked up in the DNS is stored in servers referred to as ‘authoritative servers.’” This definition is
5 consistent with the use of the term in the art. For example, further characterizations of the term “authoritative server” include:

10 There are two common types of DNS servers: the Authoritative DNS Name Server and The Non-Authoritative DNS Name Caching Server. ... An Authoritative DNS server is the authoritative source for all DNS requests made for a designated zone or domain.
 (<http://www.more.net/technical/dns/overview.html>)

The applicant is unable to find any indication that the servers taught in Hudson are DNS servers, much less authoritative servers for DNS information. The Applicant, therefore,
15 requests that the Examiner provide support for the suggestion that a person of ordinary skill in the art would interpret the term “authoritative server” as suggested by the Examiner, or allow Claim 6.

Regarding Claims 7 and 8, it is the Applicant’s position that Claims 7 and 8 are allowable for at least the reasons discussed herein with respect to Claim 1.

20 **Regarding Claims 31 and 41**, it is the Applicant’s position that Claims 31 and 41 are allowable for at least the reasons discussed herein with respect to Claims 26 and 40, from which they depend, respectively.

Regarding Claims 9, 12, 16, 18, 24, and 32, it is the Applicant’s position that Claims 9, 12, 16, 18, 24, and 32 are allowable for at least the reasons discussed herein with respect
25 to the claims from which they depend.

Regarding Claim 42,

Claim 42 recites:

42. The method of claim 40, wherein the received request for DNS information includes a request for an IP address.

5 In rejecting Claim 42, the Examiner states “Claim... [is] rejected based on the same rationale as the rejection of Claim 9.” This rationale includes “Hudson teaches that the DNS information includes an IP address (e.g. see paragraph [0003]).” The Applicant traverses this statement.

 Specifically, even if the caches of Hudson include an IP address, the Applicant is
10 unable to identify any teaching within the cited art that includes receiving a “*request for an IP address*.” Any requests taught in Hudson appear to be for “web content data” rather than IP addresses. Paragraph [0004] of Hudson characterizes “web content data” as including “web pages” and “images.” While paragraph [0003] and other sections of the cited art include the use of IP addresses to locate proximal caches, there does not
15 appear to be any teaching of a *request* for these IP addresses. These IP addresses are merely used as a tool for finding “web content data.” They are not, themselves, requested. The Applicant, therefore, requests that the Examiner more specifically point out a teaching of a “*request for an IP address*” within the cited art, or allow Claim 42.

Regarding Claims 14, 15 and 30,

20 **Claim 14 recites:**

14. The computer network of claim 13, further including means for storing data in the first cache such that a time required to examine the first cache is essentially constant as a function of a number of labels comprising the domain name.

25 **Claim 15 recites:**

15. The computer network of claim 13, further including means for storing data in the first cache such that a time required to examine the first cache is essentially constant as a function of a size of the first cache.

Claim 30 Recites:

5 30. *The method of claim 26, wherein a time required to examine the answer cache is essentially constant as a function of a number of labels comprising the domain name and essentially constant as a function of a size of the answer cache.*

In rejecting Claims 14, 15 and 30, the Examiner states:

10 Hudson teaches ... that the time required to examine the first cache is essentially constant as a function of a number of labels comprising the domain name, i.e. the first cache is the local cache (a proximal cooperative cache) and since the number of cache entries to search in this local cache is fixed/constant all the time, a time required to examine the first/local cache is essentially constant as a function of (i)
15 a number of labels comprising the domain name and (ii) a size of the first/local cache (e.g. see paragraph [0006]).

First, the Applicant is unable to identify a basis for the Examiner's statement that "the number of cache entries to search in this local cache is fixed/constant all the time," and requests that the Examiner provide support for this statement. It is the nature of web
20 content caches that the number of entries they include is variable over time. For example, FIG. 7 element 97 of Hudson teaches adding an element (web content data) to a local cache.

Second, the Applicant is unable to identify any teaching within the cited text that a search is made in a time that is essentially constant as a function of a number of labels.
25 Specifically, paragraph [0006] of Hudson does not appear to include any discussion of either a time required to examine the first/local cache being essentially constant as a function of (i) a number of labels comprising the domain name, or (ii) a size of the first/local cache. Further, it is well known in the art that searches can take a time proportional to the square or log of the number of elements to be searched.
30 While, paragraph [0003] does mention the reduction in retrieval time latency, this does not teach the limitations of Claims 14, 15 and 30. For example, there is no teaching

within Hudson that reducing retrieval time latency results in “*a time required to examine the first cache [being] essentially constant as a function of a size of the first cache,*” as recited in Claim 15. The Applicant, therefore, requests that the Examiner clarify how the teachings of Hudson relate to the limitations of Claims 14, 15 and 30, or allow Claims 14,
5 15 and 30.

Third, Claim 30 recites two different time dependencies. The first including “*wherein a time required to examine the answer cache is essentially constant as a function of a number of labels comprising the domain name,*” and the second including “*wherein a time required to examine the answer cache is essentially ... constant as a*
10 *function of a size of the answer cache.*” The Examiner does not appear to address these separate and different limitations. The Applicant, therefore, requests that the Examiner point out teachings of all the limitations of Claim 30 within the cited art, or allow Claim 30.

Regarding Claims 33 and 36,

15 **Claim 33 recites:**

33. *A method of storing data in a cache, the method comprising:*
 requesting DNS information;
 receiving data in response to the request;
 classifying the response received; and
20 *storing the data received in either a referral cache or an answer cache based on the classification.*

In reference to Claim 33, the Examiner states “Hudson teaches ... requesting DNS information; receiving data in response to the request; classifying the response
25 received; and storing the data received in either a referral cache or an answer cache (i.e., in the caching server) based on the classification (e.g., see paragraph [0063]).” The Applicant traverses this statement.

First, it is unclear to the Applicant which part of Hudson is thought by the Examiner to teach “*requesting DNS information.*” As discussed elsewhere herein (e.g., in regard to Claims 11 and 42), all of the requests taught in Hudson appear to be related to “web content data” which is characterized as web pages and images (paragraph [0004]).

5 Even if, for the sake of argument, one were to assume that the web content caches of Hudson included DNS information, there does not appear to be any teaching of requests for this information. Thus, Hudson does not teach “*requesting DNS information.*” The Applicant, therefore, requests that the Examiner specifically point out teaching of “*requesting DNS information,*” or allow Claim 33, and those claims that depend
10 therefrom.

Second, it is unclear to the Applicant which part of Hudson is thought by the Examiner to teach, “*classifying the response received.*” The Applicant is unable to identify any teaching within Hudson of classifying a received result. The Applicant, therefore, requests that the Examiner specifically point out teaching of “*classifying the
15 response received,*” or allow Claim 33, and those claims that depend therefrom.

Third, it is unclear to the Applicant which part of Hudson is thought by the Examiner to teach, “*storing the data received in either a referral cache or an answer cache based on the classification.*” As discussed in regard to Claim 1, Hudson does not teach an answer cache or a referral cache. Further, in regard to this claim element, the
20 Examiner cites storage of data in the caching server of Hudson. However, the Applicant is unable to find any teaching that this storage is based on a classification of data received in response to a request. Thus, even if one were to assume for the sake of argument that Hudson taught a referral cache and an answer cache, these limitations of Claim 33 are not

taught by Hudson. The Applicant, therefore, requests that the Examiner specifically point out teaching of “*storing the data received in either a referral cache or an answer cache based on the classification,*” or allow Claim 33, and those claims that depend therefrom.

- 5 **Regarding Claim 37**, the Applicant believes that Claim 37 is allowable for at least the same reasons as Claim 33, from which it depends.

Regarding Claim 40,

Claim 40 recites:

40. *A method of caching DNS information, the method comprising:*

- 10 *requesting DNS information;*
 receiving data in response to requesting DNS information;
 classifying the response received as an answer response or a referral response;
 storing the response received in either a referral cache or an answer cache based
 on the classification, the answer cache including a flat data structure;
15 *receiving a request for DNS information corresponding to a domain name;*
 examining the answer cache to find answer information, responsive to the
 received request; and
 responsive to the examination of the answer cache, searching the referral cache.

- 20 Regarding Claim 40, the Examiner states “Claim 40 is also rejected based on the same rationale as the rejection of claims 1 and 33.” The Applicant believes that Claim 40 is allowable for at least the same reasons as Claims 1 and 33.

- In addition, Claim 40 includes limitations not found in either Claims 1 or 33. For example, neither of these claims includes “*classifying the response received as an answer*
25 *response or a referral response.*” The Applicant is unable to find any teaching of these limitations within Hudson. The Applicant, therefore, respectfully requests that the Examiner specifically point out these limitations within the cited art, or allow Claim 40, and those claims that depend therefrom.

Claim 40 includes further limitations not found in either Claims 1 or 33. For example, neither of these claims includes “*receiving a request for DNS information corresponding to a domain name.*” The Applicant is unable to find any teaching of these limitations within Hudson. The Applicant, therefore, respectfully requests that the
5 Examiner specifically point out these limitations within the cited art, or allow Claim 40, and those claims that depend therefrom.

Rejections under 35 USC § 103

**Claim 10 is rejected under 35 U.S.C. 103(a) as being unpatentable over Hudson in
10 view of Ramanathan et al. (USPN: 6,182,136).**

Regarding Claim 10,

Claim 10 recites:

10. The caching server of claim 1, wherein the DNS information includes an MX record.

Regarding Claim 10 the Examiner states “Hudson teaches the claimed invention
15 as describe above but does not clearly teach that the DNS information includes an Mx record,” and cites Ramanathan et al. as teaching the inclusion of an Mx record in DNS information. The Applicant traverses this statement.

First, as discussed above in regard to Claim 1, Hudson does not teach all the limitations of Claim 1 from which Claim 10 depends. Thus, even in combination with
20 Ramanathan et al., not all of the limitations of Claim 10 are taught by the cited art.

Second, in the rejection of Claim 1, the Examiner appears to suggest that the “internet protocol addresses” taught in paragraph [0003] of Hudson teach the recited “*DNS information.*” These addresses of Hudson are used to forward requests for web pages and images to web caches. It is not clear how an Mx (mail exchange) record would

be useful for such a propose (i.e., forwarding a request for a web page). The Applicant, therefore, requests that the Examiner clarify how the combined art would benefit from the inclusion of an Mx record in a request for “web content data,” or allow Claim 10.

5 Third, as a motivation to combine the cited art, the Examiner appears to merely state that the combination is “advantageous.” This statement is unsupported and also made with the benefit of hindsight. The Applicant, therefore, requests that the Examiner provide a motivation to combine an Mx record with a request for web content data, from within the cited art as required for a prima facie case under 103(a), or allow Claim 10.

Applicant believes that all pending claims are allowable and respectfully requests that the Examiner issue a Notice of Allowance. Should the Examiner have questions, the Applicant's undersigned representative may be reached at the number provided.

5

Respectfully submitted,
Andreas Gustafsson

10 Date: December 16, 2005



Steven M. Colby, Ph.D. Reg. No. 50,250

Carr & Ferrell *LLP*
2200 Geng Rd.
Palo Alto, CA 94303
Phone (650) 812-3424
Fax (650) 812-3444

15